

Chapter 5

Installing and Configuring Netscape Enterprise Server and LiveWire

CONTENTS

- [Overview of the Enterprise 2.0 Server](#)
- [Enterprise 2.0 Features](#)
 - [General Features](#)
 - [Administration](#)
 - [Security and Performance](#)
- [Pre-Installation Requirements](#)
 - [Configuring Windows NT](#)
 - [Installing TCP/IP](#)
 - [Configuring Security](#)
- [Installing the Enterprise Server](#)
 - [Server Setup: Selecting Hostname](#)
 - [Administration Server Setup: Choosing Administration Access Username](#)
 - [Administration Server Setup: Choosing Administration Port Number](#)
 - [Administration Server Setup: Choosing Administration User](#)
 - [Web Server Setup: Choosing Document Root](#)
- [Configuring Enterprise Server](#)
 - [System Settings](#)
 - [Access Control](#)
 - [Encryption](#)
 - [Programs](#)
 - [Server Status](#)
 - [Configuration Styles](#)
 - [Content Management](#)
 - [Index Documents](#)
 - [Auto Catalog](#)
- [Installing and Configuring LiveWire](#)
 - [Installing LiveWire](#)
 - [Configuring Enterprise Server for LiveWire](#)
 - [Using LiveWire](#)
- [Virtual Hosts](#)
 - [Installing a Hardware Virtual Server](#)
 - [Installing a Software Virtual Server](#)
- [Netscape FastTrack Server](#)
- [Troubleshooting](#)
 - [Detecting Problems with Event Viewer](#)
 - [Monitoring with Performance Monitor](#)
 - [Online Help](#)

In this chapter, you are introduced to the Netscape Enterprise Server and LiveWire, and given step-by-step

instructions to install and configure these two applications. This chapter will go through the basic steps needed to prepare windows NT for installation of the Enterprise Server, and then downloading, installing, and configuring the Enterprise Server and LiveWire.

The Netscape Enterprise 2.0 Web Server is a third-generation server from Netscape Communications, creators of the popular Netscape Navigator software. Due to the vast number of enhancements available on this server, it is impossible to detail the operations of each and every option. Instead, this chapter will outline basic installation, briefly describe each of the available options, and provide step-by-step instructions for configuring some of the more popular features.

LiveWire is an add-on to the Netscape Enterprise Server, comprised of three main components. The Site Manager and LiveWire compiler, the LiveWire server extension, and Netscape Navigator Gold. LiveWire serves two main purposes, to help manage a Web site, and to help create dynamic content. Even though this chapter is primarily geared toward installing the Enterprise Server and LiveWire on a computer running Windows NT, the downloading and configuration information presented can be used on most platforms (including UNIX) supported by the Server.

In this chapter, you will learn the following:

- How to obtain the Enterprise Server and LiveWire
- How to configure Windows NT to the install Enterprise Server
- How to install the Enterprise Server
- How to configure the Enterprise Server
- How to install and Configure LiveWire

Overview of the Enterprise 2.0 Server

Before considering installing the Netscape Enterprise Server, it is a good idea to check the supported configurations (see table 5.1) to see whether your current hardware and operating system are capable of running the Server.

Table 5.1 Supported Configurations

Vendor	Architecture	OS	Memory Requirements
UNIX			
Digital	Alpha	OSF/1.3.2C	32 MB
HP	PA	HP-UX 9.x, 10.01	32 MB
IBM	RS/6000	AIX 3.2.5,4.1	32 MB
SGI	MIPS	IRIX 5.4, 6.2	32 MB
Sun	Sparc	SunOS 4.1.3, Solaris 2.4, 2.5	32 MB
Windows NT			
Digital	Alpha	NT 3.51, NT 4.0	32 MB
Intel	x486,Pentium	NT 3.51, NT 4.0	32 MB

NOTE

Netscape recommends 30 MB of free disk space for installation, as well as 30 MB free disk space for log files (for a server with approximately

300,000 accesses per day).

Enterprise 2.0 Features

The Netscape Enterprise 2.0 Server has all of the standard features available on most popular WWW server packages, as well as a few special enhancements that make it particularly well suited for Intranet use.

General Features

To make Web publishing easier and more convenient, the Enterprise Server can be used in conjunction with Navigator Gold, a WYSIWYG HTML editing and publishing tool. Using Enterprise's remote file manipulation feature, Navigator Gold allows users to update Web files from any remote location that is networked to the Server. The familiar Web browser interface of Navigator Gold saves the user from having to learn to use other tools such as FTP and HTML. A revision control system allows multiple users to simultaneously work on documents without risking the integrity of the files on the Server.

A fully integrated, full-text search engine allows such features as full-text and field searches, incremental indexing, and add-on support for document types such as Portable Document Format (PDF) without having to revert to third-party software. An integrated cataloging system automatically generates catalogs of files on a site based on the creating author, the creation date, or through user-defined classifications. This catalog provides users with a quick overview of all the files available on the Server.

Administration

Administration of the Server has been simplified, and the interface has been made clear and logical. The new Administration Server has several new features, including new log analysis tools that create summaries on Web site statistics such as total hits, total number of unique hosts, and total traffic transferred. The new analysis tool also supports enhanced features such as identifying which clients are accessing the most number of pages and downloading the most information.

Support for multiple domains has been vastly improved over previous versions of the Netscape Server, enabling administrators to easily host multiple sites on a single machine. LiveWire, a standard component of the server, is a visual site management tool that allows administrators to view and restructure entire sites in a graphical form. To improve remote monitoring capabilities, the Enterprise Server now supports SNMP 1 and 2, allowing administrators to monitor their servers remotely using standard SNMP capable tools.

Security and Performance

To address security concerns, Netscape has upgraded its security with enhancements such as Secure Sockets Layer (SSL) 3.0 support, advanced access control, and client-side certificates. These components help to secure not only commerce but all communications that travel through the Server.

Netscape's second incarnation of its commercial Web Server software, the Enterprise 2.0 Server, combines new caching technology, platform-specific optimization, and multiprocessor support to create one of the best performing servers available on the market today.

Pre-Installation Requirements

The installation process of the Enterprise Server comes in several steps. First, it is necessary to configure Windows NT with the proper software and network settings so that the server can be installed. Next you need to obtain all of the software that you need to get the Enterprise 2.0 Server operating on your NT

machine. Finally, you can customize your system to your needs by accessing the Administration Server.

Configuring Windows NT

The first step in preparing for the installation process is to make sure that Windows NT is properly configured. The following instructions assume that you already have Windows NT installed and running properly, and that you have already configured your network interface card.

NOTE

To install the Enterprise 2.0 Server, you must have the Windows NT 3.51 service pack #4 installed. You can find this service pack as well as other configuration tips for Windows NT Server at the NT server home page at <http://www.microsoft.com/NTServer/>. Users of Windows NT 4.0 do not need to install any of the service packs.

To successfully operate an Enterprise Server, you must have the TCP/IP protocol installed on NT, and you must have a permanent IP address assigned to your Server. You will need to know your IP address to properly install your Server and for other people to be able to reach your Server.

TIP

To help users access your system, it is helpful to register a domain name in DNS for your permanent IP address. By doing so, users can access your server using the domain name (such as ian.digiknow.com) rather than a complicated IP address (192.147.147.142 would be the IP address associated to ian.digiknow.com).

Installing TCP/IP

Before installing TCP/IP on Windows NT, you will need to know the following information:

- The IP address(es) to be assigned to your computer
- The Host name(s) corresponding to the preceding addresses
- The IP addresses of the DNS servers you will use
- The Domain name in which you will operate
 1. Open the Network Control Panel and click on the Add Software button.
 2. From the pull-down menu, highlight TCP/IP Protocol and Related Components and click on Continue.
 3. Select the components you wish to install and click on Continue (At the very least you will want to install "Connectivity Utilities" and "Simple TCP/IP Services"). You will be prompted for the location of the NT software distribution disk.

Figure 5.1 : The Network Control Panel allows you to add TCP/IP services and to configure them for use with the Enterprise 2.0 Server.

NOTE

If you want to use the Windows NT Performance Monitor to monitor TCP/IP statistics, you will have to install SNMP service. SNMP will also allow your computer to be administered remotely using remote management tools. If you choose to install SNMP, you will be prompted by the SNMP configuration dialog box. Unless you have particular needs for SNMP, you can simply select the OK button.

4. If you have installed RAS for dial-up access to the Internet (either via modem or ISDN), you will be prompted to configure RAS to support the TCP/IP. If you do, click OK.

NOTE

It is possible to use the Enterprise Server on a dial-up IP connection (either via modem or ISDN). For proper operation, you must have a static IP address defined, however. If you are using a dial-up connection to install the server, make sure the connection is live and running before installing the Enterprise Server.

5. Once the computer finishes adding TCP/IP services, click OK on the Network Settings Configuration Box. A TCP/IP Configuration box will appear allowing you to configure your network adapter.
6. From the Adapter pull-down menu, select the Adapter you wish to configure.
7. Under IP address, enter the IP address for your host.
8. Under Subnet Mask, enter the subnet mask for your host.
9. Under Default Gateway, enter the default gateway for your network.
10. If you are using DNS servers (if you are on the Internet, then you are), click DNS to bring up the DNS configuration dialog box.
11. Under Host Name, enter the host name of your server.
12. Under Domain Name, enter the domain in which your host is registered.
13. In the DNS Search Order box, enter the name(s) of the DNS servers you will be using for host lookups. Click OK to get back to the main configuration menu.
14. Now that you have TCP/IP properly configured on your NT Server, the last step is to reboot the machine to put the changes into effect. Click OK in the TCP/IP Configuration box and when prompted, restart the computer.

Configuring Security

The downside to the benefits that networking computers provides is that it also brings up several security issues. To meet this concern, Windows NT has several layers of security available to it, including user-account security and file system security using the Windows NT File System (NTFS).

Every operation that takes place under NT can be identified by the user name used to start the particular operation. The User Manager application allows you to set which resources a particular user is authorized to use on the computer, as well as which files they are allowed to access. To limit the level of access that the Enterprise Server has to your computer, it is recommended that you create a nonprivileged user account for the Server to run under. This account should be restricted to access only what is necessary to start up and operate the Enterprise Server software.

NOTE

By default, the Server uses the LocalSystem account under NT and the nobody account under UNIX. Under both systems, however, it is still recommended that administrators create a separate account for running the Server.

Installing the Enterprise Server

There are two ways to obtain the Enterprise server. The first (and the quickest) is to download the server from Netscape's home site at <http://www.netscape.com>.

TIP

Netscape allows a 60-day evaluation trial of all of their software packages, including Servers and browsers. If you want to have support, you can also purchase the server directly online and get 90 days of free technical support.

After you find the Server Download page, use the pull-down menus to select the files you wish to download. (You will need to specify the product you wish to download, the Operating System you are running, and the file type to download: .zip, .exe, or .gzip.) While you are at the Netscape site, you might also want to pick up a copy of the Navigator Gold software, which you will need to configure your Server if you do not have a Web browser installed yet. (It also allows for remote updates of Web files stored on the server.)

If you are not one for long downloads, for a negligible shipping and handing charge you can order an evaluation copy of the server on CD-ROM from Netscape's Web site at <http://www.netscape.com>. This CD will come with the Enterprise Server as well as the Navigator Gold Software.

TIP

You might want to install the Navigator Gold software before running the setup for the Server so that you will be able to jump right into the Administration Server after you are finished with the Enterprise Installation.

From the command prompt or from File Manager, run the executable file you have just downloaded (if you have the file on CD, run the setup.exe file).

TIP

Make sure that you shut down any other applications that are running before installing this software. If you are already running a WWW server, disable it in the Services Control Panel.

The setup program will first prompt you with the destination directory in which to install the Enterprise Server. If the directory you designate does not exist, don't worry; the setup program will create it for you. (Keep in mind, though, that you will need at least 30 MB of free space for the installation.)

NOTE

During setup, the server will come up with several queries. It is usually safe to use the default entries because you can change the values later using the Administration Server.

If you already have a Enterprise 1.1x server installed, the setup program will question whether you want to upgrade the server or if you want to install a new server. (You can run both concurrently, just not on the same port.) Either option will not write over your current 1.1x installation, however, and you can always reactivate it from the Services Control Panel.

Server Setup: Selecting Hostname

The Enterprise Server will automatically obtain the settings for your hostname as you entered them in your TCP/IP Configuration menu when you installed TCP/IP. If your server does not have a proper DNS entry set up, you should enter the IP address under which the server is running; otherwise, you will not be able to access your server properly.

Administration Server Setup: Choosing Administration Access Username

The administration access username is the name and password you will use to connect to and administer your Enterprise Server. Because the Enterprise Administration Server can be reached via the Netscape Navigator browser from any site on the Internet, it is crucial to set this feature for security purposes.

Figure 5.2 : If a host name does not automatically come up, it means that you have not yet properly configured TCP/IP. Without a proper hostname, you will not be able to access your server once it is installed.

Administration Server Setup: Choosing Administration Port Number

The Administration port is the port number on which you wish to operate your Administration Server (for example: <http://www.myserver.com:8888>). The installation process will randomly select a default port from the available ports on the system. You will need to remember this port so you can later be able to access your Administration Server and make changes in your Enterprise Server configuration.

CAUTION

Netscape randomly chooses a port for the Administration Server for security purposes. Because this is the gateway into your Server configuration, it would be unwise to run it on a port that could be easily identified by others (even though it does prompt you for a password).

Administration Server Setup: Choosing Administration User

This name is the name of the user under which the Administration Server (as well as the Enterprise Server) will be run. At this time, you can leave this option as is because you will be able to make changes later within the Administrative Server.

Figure 5.3 : By default, the administration server runs as LocalSystem. For security purposes, you might wish to change the name under which the Administrative Server runs.

Web Server Setup: Choosing Document Root

The document root is the highest level directory visible to the Enterprise Server. You have to specify the full path to the location where your Web documents will reside. If the directory does not exist, it will be created for you.

CAUTION

This setting is vital because a large portion of Web security is based on the premise that the server cannot access files outside of this directory tree. Make sure that the server root directory tree does not contain any files that should not be accessible by the guests to the system.

Click finish to complete the installation and to start up your Server. If you have already installed the Navigator Gold browser, it will start itself up and display the new default home page on the Enterprise Server.

TIP

If at this time, the Server does not start up, you can use the Event Viewer in the Administrative Tools folder of Windows NT to see what caused the Server to fail.

Figure 5.4 : If the installation is complete, and TCP/IP is configured properly on your system, the Enterprise Server Home Page should appear within Navigator Gold.

Configuring Enterprise Server

To start configuring your new Enterprise Server, open your Web browser to the following URL:

`http://www.yourmachine.com:<admin port>`

TIP

Forgot what the administration port is? (You were warned to remember it.) No problem. You can go to the admserv\directory in your installation directory and look at the contents of the ns-admin.conf file with a text editor. This file contains information on which port is running the Administration Server.

When you first connect to the Administration Server, you will be prompted for the login name and the password for your Administration Server. After you enter these, you will be welcomed by the Enterprise 2.0 configuration page.

To view and/or change the configuration of your server, click the server name. Doing so will bring up two menu bars in the Server Configuration page. The bar across the top has the main headings for configuration. The tool bar in the left frame shows the individual settings available under each main heading.

Figure 5.5 : When you are entering the Administation Server you will be prompted for the user name and password which you entered during installation.

Figure 5.6 : Once you successfully enter the Administrative user name and password, you will see the Enterprise 2.0 Configuration Page. From here you will be able to administer all the servers installed on your system.

There are too many options available on the Server Configuration page to outline them all in this chapter. Instead, you will be provided with a general overview of what each "section" of configuration options does, and how it can be utilized in an Intranet setting.

System Settings

The system's settings are the critical configuration options that determine how your Server operates. Settings include turning your server on and off, setting default values for server configurations (see table 5.2). At any time, you can use the restore configuration to restore previous settings of your server. (A great feature when you first start experimenting with the configuration files and all of a sudden, everything stops working.) By default, the Server will store the last 10 changes. To increase this number, click Configure Backups and enter the number of backups to store in the Number of backups field.

Figure 5.7 : Once you select a Server from the Configuration screen, you will be given a variety of options for modifying settings on a particular Server. View Server Settings.

This option provides a quick overview of the following Server variables:

Setting	Default Value	What It Does
Server Root		This is the directory in which your server files will be installed.

Hostname	Name of your computer	This is the name by which your server can be reached by others.
Port	80	This is the default HTTP Port under which your server can be accessed.
Error Log	ROOT\logs\errors	This is the file that stores the errors that the Enterprise Server detects.
DNS	off	This determines whether or not the server will execute a domain name lookup on hosts accessing the server.
Security	off	This determines whether your server is capable of secure transactions.
Additional ns-icons Documents mc-icons Directory		These are directory trees outside of the document root that can be accessed by the server.
Primary Document Directory		This is the main document tree that is accessible to the server and houses HTML files.
Index Filenames	index.html home.html	These are the files the server will look for if user does not specify a file in a directory.
Default MIME type	text/plain	This defines the default method that files are sent by the server if the server can't determine the proper type for the document.
Directory indexing	fancy	This option sets how the server responds if it cannot find one of the default file types in a directory.
Access log	ROOT/logs/access	This is the location where the Server stores a record of the hosts accessing files on the Server.

NOTE

Making changes in fields is a two-step process. The first step is to enter the changes and click OK. Next, you will be given the option to Undo the changes you have just made or to Save and Apply the changes. Clicking Save and Apply will make the changes and will restart your server to put the changes into effect.

Performance Tuning

This option allows you to configure the DNS settings for the Server. Turning DNS on will cause the Enterprise Server to look up the domain name of every computer that accesses files on the site. While this does improve security and tracking capabilities, it also causes a heavy load penalty on high-traffic systems. To help alleviate some of this load, the Server has a new feature that allows you to cache DNS entries. After performing an initial lookup, the Server then stores the results so that subsequent accesses by the same host do not have to be looked up.

You can specify the size of the cache as well as the amount of time before an entry is expired. By default, the Server will cache 1024 entries and expire them after 1200 seconds (20 minutes). You can specify

between 32 to 32768 entries, and set an expiration time between 1 second and 1 year (specified in seconds).

CAUTION

While disabling DNS can lead to performance gains on busy systems, it disables hostname restrictions, and hostnames will not appear in the log files. Log files will contain IP addresses instead, and host-based access restrictions have to be based on IP numbers rather than domain names.

Network Settings

This option allows you to set the name of the user under which you wish to run the Enterprise Server. To change the user name, you will need to enter the new name under which you wish to run the Server, as well as the password set for this user. Without the proper password, your Server will be unable to restart properly. If you ever change the hostname of your server, or if you set up an alias for the server, you would register this change in the Server Name field. The Server Port field sets the port that the server listens to for incoming requests. By default, the HTTP port is 80, and the standard HTTPS port is 443. Technically, the port number can be any port from 1 to 65535. Bind To Address: shows you the current IP address that the server is responding to.

CAUTION

Changing to a nonstandard port means that users will be required to add the port number to the URL they use to access your site. For example, if you use port 8080 for your server, your URL is

`http://www.yourhost.com:8080.`

Error Responses

This option allows you to customize the message that the Server displays to a client when it encounters an error. The new error message can be either a file on the Server or a CGI script. You can set customized error messages for the following:

- Unauthorized Access
- Forbidden Access
- File Not Found
- Server Error

NOTE

For several settings, Error Responses being one of them, the Enterprise server allows you to select which files are affected by the modifications that you make.

Choose Entire Server applies your changes to every document that the server maintains.

Browse files allows you to specify files or directories to which you want to apply or deny the changes.

Choose Wildcard Patterns lets you apply your changes to files or directories you specify with wildcard patterns. This is an easy way to

specify a large number of files in separate directories (such as *.html) or files in specific subdirectories (such as /image/*).

Dynamic Configuration Files

Even though Web Server Content is often maintained by several people, for security purposes it is ill-advised to give every user access to the Administration Server. At the same time, a Webmaster could easily get inundated with requests if he has to make every small configuration change on a Server for all of the users. The Dynamic Configuration Files option allows a Webmaster to give users access to a subset of configuration options, so that they can control only those elements which they need to.

Individual users can use a configuration file called .nsconfig in their personal directories to set a number of parameters, including custom error messages, defining file types or encoding, as well as activating access control.

Access Control

Access Control serves two major purposes: managing user databases, and controlling access to directories and files on the Server. Information on individual users or groups is stored in the Servers' built-in user databases. This information can be used in a variety of ways, the most important of which is user authentication.

Since Access Control allows you to set both read and write permissions on specific directories and files, it can be used in two separate ways. The first is to restrict individuals not found in the database from accessing information on the site. (Alternatively you can also deny access to particular hosts or domains.) Since Access Control also sets write access, it can also be used to select which users are able to use Navigator Gold to remotely create or update files on the server.

Encryption

The basic concept that allows the Internet to exist, the free passing of information from one computer to another over public networks, is also its key drawback. Since information being sent from one host to another is passing over public networks, it is possible that the information could be intercepted by others. What is worse is that due to the basic fundamentals of Internet Networking, it will always be possible for users to intercept data in transit.

Does this mean that you are exposing yourself to all sorts of risks? Thankfully, no. Consider most network communications to be as safe as voice communications over the phone or sending a letter via the mail. If, however, you do have information that is sensitive to prying eyes, encryption is the solution. *Encryption*, similar to encryption on your PC, disguises information before it is sent over the Internet, making it meaningless if someone intercepts it.

To address the security issue, Netscape products use the Secure Sockets Layer (SSL). SSL guarantees that information sent over a network can not be deciphered even if it is intercepted. It also ensures the integrity of information, not allowing users to intercept, change, and resend pieces of information. Finally, it can also be used to authenticate that a piece of information was actually created by the party claiming to have created it.

Using SSL

The basic operation of SSL is simple. The Enterprise Server has a "privacy key" attached to it. One part of this key is public, the other is private. When a browser wants to send encrypted information to a server, it reads in the public part of the Server's key. It then encrypts the message with the public side of the key and

sends the encrypted message to the Server. Since only the private side of this particular key can decrypt the message, the Server is the only place where the message can be decrypted. Setting up SSL on the Enterprise Server takes several steps and involves a few different parties.

1. First, generate a Key using the "Generate Key" option. This part is the public key and will be used to create your private key.

CAUTION

If you forget your password, you will have to generate a new key-pair file and obtain another certificate (leading to additional costs). This will lead to your secure server being down until you receive the new certificate, and you will be unable to read already encrypted messages.

2. Next, request a Certificate from a Certification Authority (CA) by filling in the request form and completing the appropriate paperwork. (There is a setup fee and an ongoing fee for Server Certificates. Rates vary among Certification Authorities.)
Obtaining a certificate can take anywhere from two days to two months so plan ahead!
3. Once the Certificate comes back from the CA, you can install it on the server and activate security. Since it is possible to have multiple certificates, and since certificates run out every year, the certificate management option becomes important to ensure that all the right certificates are installed and activated.

Programs

The Enterprise Server is capable of much more than just serving up HTML, text, and graphics files. It is possible to run programs either on the client or on the Server that allow for any type of interaction you can imagine. You might allow users to use a search program on your server to find information that they need; a user might use a custom group-scheduling program to schedule a meeting; or a user might be logging her timesheets through a Java application.

The Enterprise Server currently supports three types of programs: CGI, Java, and JavaScript. CGI (Common Gateway Interface) programs can be written in any number of languages such as Perl, C, and C++. The common feature between CGI programs is that they have a standard method in which they accept and return information. Java is a full-featured programming language that was created by Sun Microsystems for use on the Internet. JavaScript is a simpler scripting language based on Java, especially useful for creating simpler Web applications.

CGI

CGI scripts can be run on the Server in one of two ways. The first is to set the server up with a cgi file type (with a .exe or .cgi extension), so that CGI scripts can be run from anywhere on the server. The other option is to specify a directory as the cgi directory and only allow files within that particular directory tree to be executed. To set up a cgi directory:

1. Under URL Prefix, enter the trailing part of the URL that indicates you are accessing a CGI script (this is usually *cgi-bin*).
2. Under CGI Directory, enter the full path to your CGI directory (such as c:\Server\cgi-bin).

CAUTION

You can use both the CGI Directory and the CGI extension settings concurrently. For security purposes and due to the nature of CGI scripts,

however, it is recommended that you keep all of the CGI scripts in one directory and that you not let inexperienced users install their own CGI scripts.

Java and Javascript

This option allows you to activate the Server's Java interpreter and to specify the Java applet directory. Similar to the CGI directory, this allows you to specify a directory on the server where all Java applets will be stored.

NOTE

No, this feature does not turn your computer into a coffee grinder. Java is a new programming language based on C++ that was designed to run as an interpreted language. Rather than have users download precompiled applications, by using Java a client can download a section of source code, interpret it, and run it on the client machine. Java applets can be used to generate on-the-fly graphic reports, query databases, conduct interactive training and provide continuous updating information.

Server Status

The most important aspects of running a server are to be able to identify which files are being accessed and how many people are accessing those files. The Access Log stores information on traffic to the server that can later be analyzed using built-in features. The Enterprise Server also allows you to monitor the Servers' usage so that you can keep it operating at its highest efficiency.

View Error Logs

Viewing the Error log allows you to keep track of any errors that the Enterprise Server encounters. If CGI scripts are failing, the error log will often tell you at which point a script might be failing. The error log will also point out when files are requested that do not exist. Often this is due to links becoming outdated or being misspelled. Correcting errors such as these will help avoid aggravating users with a "File Not Found" error message.

Monitor Current Activity

This option displays the current output of the server as well as the number of active processes the server is handling. If you are running a high-traffic server, this monitor will help determine where bottlenecks in the system might lie, and what can be done to enhance performance. (It could show that your system needs more memory to run more processes, or that the system is running faster than the Ethernet card can handle.)

Log Preferences

This option allows you to customize the location of your log files as well as the format to use. By default, the Enterprise Server uses the common log file format. You can also set whether the server will register hosts based on their domain name or IP address (if you have DNS turned off, it will always register them under their IP address). Additional logging features include the following:

- **Referer Headers:** This will log the last page that a user visited before coming to your Web site. This will show what part of your site they were visiting that led them to another area. (Shows cause and effect.)

- **User-Agent Headers:** This will log the type of browser that users are using to access your server. When making decisions on what enhancements (advanced HTML, Java, plugins, ActiveX) to use, knowing what the browser base is will help to make decisions that will benefit the largest numbers of users.
- **Query string of the URI:** This will log the entries after the question mark in a URI. Typically, this is the query string used in searches, allowing you to see what terms people used to find your information.

Generate Log

Rather than having to decipher a list of thousands of individual entries of visitors to a site, this option creates a summary of the log file entries that actually makes sense. Summary results can be displayed on-screen or saved to a file for future reference and comparison.

This option allows you to create a summary of your log file entries that actually makes sense (rather than a list of thousands of entries). By default, the summary results will be displayed to the screen but can be configured to save to a file (in HTML or text format).

TIP

Log files are often overlooked as not being vital. The log files, however, are a direct evaluation of your Web server, showing what areas are popular and what areas are not. They also give vital statistics on the number of hosts connecting to the server, as well as the amount of traffic that different areas of the server receives.

After generating a summary report of your Access Log, you might wonder what the numbers mean. Here is a summary of some of the important figures:

- **Total hits:** This count tells you how many pieces of information were downloaded from your Server. A single page accessed by a client containing 20 small graphics would be registered as 21 hits. Because of the way the counter works, this number has little value, other than it sounds nice being able to claim that your server gets 100,000 hits a day.
- **Total unique client hosts:** This count gives the closest estimate of the number of individuals that are accessing servers. The true meaning of this number is a count of the number of individual IP addresses that are accessing your system. However, it is possible that more than one person might be using the same IP address, so the number of total users lies somewhere above the number of unique hosts.
- **Total kilobytes transferred:** This count shows how much outgoing traffic is leaving your server. This might be a vital number if you are paying for bandwidth on a usage basis.
- **Top X periods:** This shows the times of the day that your server is the busiest.
- **Most commonly accessed URLs:** This shows the individual directories and files that are accessed most often.

Configuration Styles

Configuration styles are a quick and simple way to apply a set of configuration options to specific files or directories on your Server. You could for example set up a configuration style that configures how to handle access logging, how to handle errors, and where to look for cgi scripts. The style can then be applied to files or entire directories, saving the time of having to individually set options for files and directories.

Content Management

Because it is likely that your computer will be used for purposes other than as a Web Server, you probably want to be able to limit the areas that the Web Server can access. The Enterprise Server allows you to

configure which areas of your server are accessible to the public, as well as which files are sent by default to a client.

The Primary Document Directory is the highest level which is readable by the web Server. Optionally, however, you can set Additional Document Directories which map to other directories outside of the Server document tree. You could, for example, have the directory /images/ point to d:\images, even though the server root is c:\Server\docs. This option is especially helpful if you are running short on drive space and want to move high disk volume directory trees to another drive or partition.

Remote File Manipulation

One of the most convenient features of the Netscape Enterprise Server is that it allows users to update files remotely via the Navigator Gold software. Using the familiar Navigator interface, combined with a simple WYSIWYG HTML editor, updating files on the Server has become quick and easy.

CAUTION

If you turn on Remote File Manipulation, you should use Access Control (as mentioned in "Configuring Netscape Server" and "Access Control") to restrict the users that are allowed to write to particular files or directories. Otherwise, anybody will be able to edit your files.

Document Preference

This option allows you to set the files that the Server looks for when a user enters an URL that ends in a directory name. By default, it will look for a file called index.html and then a file called home.html in a directory. To change or add to this list of files, under Index Filenames: enter the names (in order) of the files that you want the Enterprise Server to look for. (Use commas to separate file names.)

If one of the default files is not found, by default the Server will show a directory listing of the files available in the directory, along with graphics to depict the type of each link (directory, file, sound, image, and so on). To turn off graphics, set directory indexing to simple. To turn off directory listings all together, set directory indexing to None.

NOTE

Turning off directory indexing is often a good way to protect your files. If you have a directory containing only images, a user could easily back into this directory and get a full listing of all your graphics. While theoretically they have access to each of them through the separate Web pages that incorporate these images, there is no reason to make it easy for someone to grab all of your graphics.

In certain situations, you will want your home page to be something other than one of the default file names, or a file other than one in the root directory. If so, click home page and enter the name of the file you want to set as your default home page.

URL Forwarding

As Websites develop, files are bound to move around the server, often onto another server. URL Forwarding allows you to set up automatic forwarding for particular files and directories which have been moved elsewhere.

Virtual Servers

A common need is to have one computer act as if it is hosting multiple web sites. Both Hardware Virtual Servers and Software Virtual Servers allow you to create the appearance that you are running multiple Servers. (See the section "Virtual Hosts" for more details.)

Document Footer

This option allows you to specify the text for a generic footer to add to the selected files. This is especially handy when you want to install navigation bars in the bottom of pages that may change in the future. (With the footers, you have to edit only one file to make changes on all the pages, rather than having to edit all the affected files.)

Version Control

Enabling Version Control prevents two users from editing a single file simultaneously. If Version Control is activated, users can "check-out" documents, preventing others from editing the file. (Other people can still see the document, they just cannot modify it.) Only once a person "checks-in" a document will it be available for others to edit again.

Index Documents

Up until now, searches on Web sites required the installation of third-party search software packages such as Glimpse and WAIS or CGI scripts to do simple searching. The Enterprise Server comes with a built-in, full-text search engine that can quickly be configured to offer your users a quick and easy overview of files on your Server. The search engine allows you to create "collections" of pages and directories, which can then be searched for key words.

Collections are easy to create, and once created can be automatically recreated at specified intervals, or can be edited manually at any time. To create a new searchable collection:

1. Under Collection Name:, enter the name of the collection.
2. Under Description:, enter a brief description of the collection.
3. Under directory to index:, enter the full path of the files to be indexed in this database.
4. Specify whether you want to include subdirectories.
5. Specify the file types to index.

Auto Catalog

Often, a search engine by itself does not provide users with the information they are looking for, nor does it supply them with other relevant information. To address this need, the Enterprise Server comes with a Cataloging system that catalogs files on the server based on modification date, title, author, and a user-defined classification. This allows users to quickly access all files created by a certain author or to view all the files that have been updated or created in the last few days.

Similar to the searchable collections, Catalogs are easy to create, and once created can be updated automatically at specified intervals or can be updated manually at any time.

Figure 5.8 : Creating a new collection is as simple as naming it and entering the path of the files to be included in the database.

Installing and Configuring LiveWire

LiveWire is an add-on for the Netscape Enterprise Server which allows administrators to create client-server applications that run over the Internet. Using JavaScript, a variety of programs can be written to create dynamic HTML pages that process user input and maintain data both in files and relational databases. Applications could include in-house on-line training sessions with interactive tests, Intranet publishing, order tracking, or even something as simple as timesheets.

LiveWire comes in three parts, the Site Manager and the LiveWire compiler, the LiveWire server extension, and Netscape Navigator Gold. LiveWire Pro, which comes bundled with an Enterprise Server 2.0 purchase, also comes with a Structured Query Language (SQL) database and report generator.

Installing LiveWire

At this point, you should already have installed the Netscape Enterprise Server, as well as Netscape Navigator Gold.

1. Open the Services Control Panel and highlight your Netscape Enterprise Server. Click the Stop button to halt the Server.
2. Run the LiveWire executable file which you have downloaded.
3. When prompted to Select HTTP Servers to Configure, select the Servers for which you wish to configure LiveWire. (In most cases you will only be running one server.)
4. When prompted to Enter Information, enter the host name under which your server is operating.
5. When prompted to Enter Information for a second time, leave the option blank and click "next."
6. Go back to the Services Control Panel, highlight your Web server and click "start."

Figure 5.9 : The LiveWire application manager is the key to creating Dynamic Web pages.

Configuring Enterprise Server for LiveWire

1. Connect to your Netscape Administrative Server through Navigator Gold.
2. Within the Server configuration, click Programs.
3. Under Programs select LiveWire.
4. Toggle "Activate the LiveWire application environment" to yes and click OK.
5. Click "Save and Apply" changes to activate LiveWire.

Figure 5.10: Once LiveWire is installed, it has to be activated from within the Enterprise 2.0 Configuration menu.

Using LiveWire

Once LiveWire is installed, it can be used in one of several ways to help enhance a site. For the Web Site novice, the Site Manager is an easy way to get started building a Web Site. Using a few standard queries, Site Manager will create an entire website based on one of the many templates included with the software. These sites can then be customized to better suit the user.

The most exciting (though most difficult to master) feature of LiveWire is that it helps users develop client-server applications. These applications can serve any number of purposes, from simple mathematical calculations to complete database management.

Virtual Hosts

One of the new integrated features in the Enterprise Server is that it allows you to run multiple Web sites on the same machine. This might be helpful when one department wants to have a Server accounting.company.com, while another wants to have a Server legal.company.com. Rather than purchase two separate Web Servers (which would get expensive if you have many departments), it is possible to run

The Enterprise Server allows two ways of running additional Servers. The first option is to install a new Server and to run it on a different port. The other option is to run multiple Servers on the same port, 80. Because port 80 is the standard port, the second option is the preferable way of accomplishing this.

There are two ways to run multiple Servers under port 80. The first is to run hardware Virtual Servers; the second is to run software virtual servers.

Installing a Hardware Virtual Server

Before making any changes to the Server, you will need to load additional IP addresses that your NT server will respond to. Under NT, follow these steps:

1. Open the Network Control Panel and double-click TCP/IP protocol in the Installed Network Software box. This will bring up the TCP/IP Configuration box.
2. Click Advanced to bring up the Advanced Microsoft TCP/IP Configuration box.
Figure 5.11: NT's Advanced TCP/IP Configuration.
3. To add an IP address to your computer, enter the IP address and the corresponding Subnet Mask in the appropriate fields and click Add.
4. Click OK to return to the TCP/IP Configuration menu, then click OK again to return to the Network Settings menu.
5. To put your changes into effect, click OK in the Network Settings box and restart NT when prompted to do so.

NOTE

By using this method, you can add only five IP addresses to your NT server, even though the Enterprise server is capable of supporting up to 16 addresses. If you wish to add additional IP addresses, you will have to edit the Windows NT system registry directly. This is an option recommended only to experienced users, as mistakes in editing the registry can cause your system to fail. (Still have that rescue disk handy?)
For more information on the subject, see
<http://www.lancomp.com/MultipleDomains/>.

6. Under Content Mgmt, go to the Hardware Virtual Servers Setting.
7. In the IP Address field, insert the IP address you just added.
8. In the document root field, insert the document root for the new Server you are installing.
9. Click OK, and then click Save and Apply to make the changes and restart the Server. Now try to open your browser to the new Server you have just created.

Installing a Software Virtual Server

Software virtual servers behave slightly differently than hardware virtual servers in that they do not require a separate IP address. Rather, software virtual servers look at the domain name asked for on the incoming request and will serve up a file appropriately. To install a software virtual server, follow these steps:

1. Under Content Mgmt, go to the Software Virtual Server setting.
2. In the URL Host field, enter the host name to which you want the Server to reply.
3. In the Home Page field, enter the path to the home page to use for the virtual server. (Typing a full path will use the specific document; typing a partial path will be interpreted as being relative to the document root set in the Primary Documents Directory setting.)

4. Click OK, and then click Save and Apply to make the changes and restart the Server. Now try to open your browser to the new Server you have just created.

NOTE

For software virtual servers to work, the host name specified in the URL Host field has to have a DNS entry pointing it to the IP address of the server.

Netscape FastTrack Server

The only drawback at this time of the Enterprise Server is the high costs associated with the server. For the business that does not need all of the features that are offered with the Enterprise Server, Netscape offers the FastTrack Server. Smaller, quicker, and easier to install, the FastTrack Server is ideal for businesses setting up their first Web Server.

FastTrack still allows users to update files remotely via Navigator Gold, supports SSL security, and comes bundled with LiveWire. Features that are lacking from the Enterprise Server include the integrated text search, revision control, the Cataloging system, SNMP support and the LiveWire Pro Database.

Troubleshooting

If nothing seems to want to work properly, don't worry. There are several options available for help when you are having trouble with your new Enterprise Server installation. NT itself comes with two management tools—the Event Viewer and the Performance Monitor—that can help to detect errors and optimize a system.

Detecting Problems with Event Viewer

Event Viewer, an application in the Administrative Tools program group, keeps a record of critical events and system errors that might give clues as to why certain operations might be failing. (For example, it would show that the Enterprise Server ran out of memory if you tried to run it on a system with only 16 MB of RAM). If you run into a situation where the server is not loading properly and the Server's Error Log does not offer an explanation, it is highly likely that the Event Viewer will have a detailed description of why the Server failed to work.

Monitoring with Performance Monitor

The Performance Monitor, also in the Administrative Tools program group, allows you to measure a handful of different performance elements. You can measure memory load, CPU load, as well as I/O information to determine where possible bottlenecks on the server might lie. If you installed SNMP when you configured your TCP/IP settings, you will also be able to monitor several TCP/IP elements.

Figure 5.12: The Windows NT Event Viewer can help lend clues as to why a Server might be failing.

Online Help

If you have tried everything, even having gone so far as to read the manuals and on-line help that come with the Server (which, by the way, are extremely helpful), don't give up yet. There are a number of online resources that can be extremely beneficial in solving problems.

Usenet Newsgroups are a great source of information on any problems you might be having. Following are some related groups:

- **comp.infosystems.www.servers.unix** is a great source of information on UNIX servers.
- **comp.infosystems.www.servers.ms-windows** is a great source for information on Windows 95 and Windows NT-based servers.
- **comp.os.ms-windows.*** offers several groups on various Windows-related issues.

Netscape Communications Corporation has several Server-related help areas available on its home site at <http://www.netscape.com>. Netscape NUGgies, or Netscape User Groups, are a group of dedicated, secure newsgroups run by Netscape. Discussion groups range from Browser discussion groups to the all-important (for you) Netscape Server User Group. NUGgies can be reached at http://www.netscape.com/commun/netscape_user_groups.html.

The Netscape Server Support page at <http://www.netscape.com/assist/support/server> has links to a wealth of server-related information, including a FAQ, an online installation guide, Technical Notes, Patches, and, if all else fails, a Help Request Form to the friendly techies at Netscape.

No matter what happens, don't give up! Web technology is still fairly new and is evolving every day. New techniques and services are constantly being added, with the result that actual documentation is often hard to find. Most importantly, don't be afraid to ask for help when you need it. There are plenty of friendly WebMasters out there that remember what it was like when they first started administering a Web site.



Chapter 23

Intranet Security

CONTENTS

- Security and Access Policies
 - Writing Policies
 - Sample Policies
 - Policy Pitfalls
 - User Education
 - Accountability
- Security Through Obscurity
 - Using Non-Standard Ports
 - Using Hard to Guess Names
 - Hiding Your Server's Name
- Using the Server Security
 - Restricting by IP Address
 - Usernames and Passwords
 - Configuring Basic Authentication
 - Other Access Methods
- Firewalls
 - Network-Level Firewalls
 - Application-Level Firewalls
- General Security
 - Modems
 - Other Remote Networks
 - Physical Security

Intranet security means not only keeping unauthorized users from reading, changing and deleting company information. It also means keeping employees from wasting time on the Internet and also from making the company look bad by posting sensitive or inappropriate information.

There are, fortunately, some things that can be done to improve the security of your Intranet. These are both technical solutions and policy decisions that can be put in place to help make things safer and more productive.

In this chapter, we will learn:

- How to write a security and access policy
- How to enforce policy decisions
- How to install software and hardware to limit exposure from the outside
- How to keep internal users from abusing the Internet
- How to spot something wrong

Security and Access Policies

One of the most important-and most overlooked-part of implementing network security is to have a written policy. A policy is a document that describes what you are trying to protect and what you are trying to protect it against, as well as describing what should and shouldn't be done over the network. Having a policy allows you to measure how effective any procedures you have implemented are. Policies also help explain to users and management why some things need to be done differently or not at all.

Writing policies is not easy and this book will not go into detail on how to do it. We will discuss generally what needs to go into a policy and discuss where to get some sample policies.

Writing Policies

Writing policies is a complicated task because it requires keeping ideas general enough to be flexible but specific enough to be useful. There are normally two different parts to the policy. The first part is the security portion, which describes what needs to be protected and how to protect it. The second part is the access part. This part describes what can be done over the network.

The first step in writing the security portion is to define what needs to be protected. This is the hardest part to determine because most users don't know what they use or how important it is. It is often better to work with each group, determine what they use for files and programs, and compile a list of who uses what.

CAUTION

When writing the policy, you should not use specific filenames because they may change. Instead it should be kept general, such as a statement reading "Payroll files should only be viewed or changed by payroll employees." Keeping the model general allows flexibility and changes to occur without having to rewrite the document.

Next you need to determine who shouldn't be able to get to what. Often security is implemented in layers or tiers. The top tier is for very sensitive machines such as those that create accounts or that enforce security. This could be the NT domain controller, NIS master, or Kerberos master. The NT domain controller handles security and accounts for NT based networks, while the NIS master and the Kerberos master handle accounts for mainly UNIX-based accounts.

The second tier of information should be allowed to be changed by only certain groups. It may need to be viewed by more than one group. Define what those groups are and what access they need.

The third tier is usually readable by everyone and writable by only certain groups. The last tier is information that can be written or read by everyone. The Internet usually falls under one of these tiers. Some companies may only allow certain people or groups of people to post information. This helps to eliminate inaccurate information from being sent out from the company.

NOTE

This security model does not specifically deal with the accesses from the Internet. It is assumed that Internet users can get access to the network, and any document that is readable by everyone can be read by these users as well. Hopefully this is not true, but if you assume it is, then you can prepare for it.

Once the part of the policy describing who can change what files is determined, usage is usually discussed. An example might be a section determining that passwords are required, must be hard to guess, and also that passwords should not be given to anyone or written down. It is common to refer to other documents that explain password procedures or other site specific requirements. This allows specific requirements to be

documented and changed without having to rewrite the security policy.

Once the access policy is determined, it is common to put in a section describing what will happen if it is not adhered to. This is often a touchy issue and it will be necessary to get the personnel department or manager involved as well as legal assistance.

Once the policy is written, it should be reviewed by either a lawyer or a legal department. Upper management needs to be informed of what the policy means and how it needs to be enforced. If upper management doesn't understand or agree with the security policy, it is not useful.

Sample Policies

Policies are general guidelines, and it is common for a template policy from another company or institution to be used as a guideline. Sample policies can be downloaded from the following sites:

- <http://delphi.colorado.edu/~pubs/draft9.html>
- <gopher://gopher.eff.org/11/CAF/policies>
- <http://www.crmwc.com/aup.htm>
- <http://chico.rice.edu/armadillo/acceptable.html>
- <http://all.net/books/policy/top.html>

Again these policies are guidelines and may not describe the requirements of your specific site. They should help explain how the document should be worded and what should be included.

Policy Pitfalls

When writing policies it is important to use careful wording and expressions. In case of disputes, this policy may need to be taken into a court and it should be reviewed by a legal department or lawyer who understands computer laws.

One of the important things to remember when writing either this policy or the security policy is to be consistent in everything you allow or deny. For example, if the access document says no access to online stock quotes, and one of the upper managers has access, then your policy is no good. There should be a way for someone to get access, even if it will not be allowed. A good way to do this is to have a clause that says "No access to online stock quotes without permissions from ____". This allows some people to have it and not others.

Employees' feelings must also be taken into account, especially in companies where employee morale is important. Improper or harsh wording can cause employees to feel "Big Brother" is watching and may get bad feelings about the company. This is not to say that employees should be allowed to do whatever they feel like on the Internet, but a little trust can go a long way.

User Education

One of the most important steps in allowing employees to access the Internet or Intranet is to teach them how to use it. There should be a document or required training class that employees must take before being unleashed on the Internet. This class or document needs to cover a few issues. Among them are:

- Netiquette. This is like etiquette in the network world. This section of the document should discuss the simple issues like not posting blatant advertisements to groups that don't allow them, the proper way to subscribe and unsubscribe from mailing lists, and not forwarding personal e-mail to mailing lists. There is a frequent posting in news.answers that discusses netiquette. This posting is a good reference to use when developing a user education program.

- How the software works. This is simply sitting down and teaching users what hypertext is, how it works, what is available and how to find that which is not readily available. This section should cover any Internet tools the company uses including News readers, WWW browsers, FTP clients, and E-mail.
- Security. This includes basic security guidelines including e-mail insecurities, choosing good passwords, and not sending company sensitive documents to mailing lists.
- Accountability. Explaining to users that connections can be logged is a way to reduce the amount of time wasted on the Internet. This is only possible in certain network configurations where the access is through a particular point, such as with InterNotes, Socks, or proxy servers.

If the document, or class, covers these basics and teaches the users how to act responsibly on the Internet, many problems can be avoided.

A good example to show employees that what they post is reflected back on the company is to search one or two of the major search engines for the company name and see what comes back. Everyone might be surprised!

Accountability

In our list of topics users need to be educated on, we mentioned accountability. This can be used to help deter employees from spending too much time on personal projects.

Many companies have logs files that contain statistics on user accesses. These may include:

- Size. This is the number of bytes downloaded per connection. This statistic allows departments or users to be billed, or help account for bandwidth usage.
- Name. This is the document name, such as index.html or button.gif. This name might give clues as to what the document contains.
- Time and date. This tells when the file was downloaded. The policy might allow personal use after work hours or during lunch.
- Site name. This is the remote site connected to. This might also give clues as to whether it is work related or not.
- Destination port. This may be FTP, HTTP or any other port. This usually gives a clue as to what protocol was used to download the file.

CAUTION

Filenames and destination ports may be misleading. Files can be called anything and do not have to end in particular extensions if the server is configured properly. Servers can also be configured to run at non-standard ports

Many sites institute a billing procedure to hold departments accountable for supplies, and consider billing for network bandwidth. This will at least have the effect of allowing management to see who is using the bandwidth.

Other sites have a "Top Ten Users" or "Top Ten Sites" list. Employees don't want their managers to see their names on the top ten users connecting to www.wastetime.com, and may limit their usage.

Using accountability to convince employees not to abuse the Internet is usually better than trying to use technical tricks to block certain accesses. It often allows for better employee relations since the company is showing that they trust their employees not to waste time, and are only trying to keep track of where the bandwidth is being used.

Security Through Obscurity

One of the easiest ways to prevent casual users from accessing your Intranet is to make it hard to find. This will not prevent someone who is trying to break into your site, but it may keep out some people.

Security through obscurity isn't real security. It is more like camouflage and, like camouflage, once your site is seen it is extremely vulnerable. Obscurity is a good way to make intruders spend time poking around and hopefully do something to set off alarms and alert someone that they are trying to get in. The best defense, though, is to use a combination of security measures, such as server security or firewalls, in conjunction with obscurity. Hopefully this will alert you to intruders knocking at the door before they can get in.

There are a few ways to hide your Intranet. The following ways are covered in detail in the next few sections:

- Using non-standard ports. The standard port is 80. Using a different port will make it harder to find.
- Using hard to guess names. Most companies use WWW for the Web server machine name. Using something different can make it harder to find.
- Hiding your server's name. This can be done by not listing it in the DNS tables for your site, and not using it to browse the Web, send e-mail, or post to Usenet.

Using Non-Standard Ports

Using non-standard ports is one of the easiest ways to hide your site from prying eyes. Most Web servers have a simple configuration that allows the administrator to change which port to listen on. The default port is 80. Changing the port from 80 to something else will make it harder to find.

Changing the Port Number with Apache

Changing the default port is fairly easy to do. For Apache, there are two ways to do this-depending on how it is run. If you run Apache via inetd, then you simply need to edit /etc/inetd.conf and change the port to something other than 80. Then restart inetd by sending it a HUP signal, for example kill-HUP pid (pid stands for *Process ID number*).

If you run Apache from a startup script such as rc.local, you need to edit the httpd.conf file, which is usually found in the httpd/conf directory. There is a line in this file that contains a port directive. By default it uses port 80. Simply change this to a different port and restart HTTPD.

Changing the Port with Netware Web Server

Changing the port for Netware Web server can be done using the Web server manager program. Simply choose file and select server. Type in or select the drive that has the Web server tree mapped to it. Change the TCP port number to something other than 80. Next click OK, and then click save and restart. Enter the Web server password when prompted and click OK.

Changing the Port with Netscape Server

To change the port in Netscape enterprise server, you need to use the server manager program. You can start this process by running start-admin and then pointing a forms capable browser at <http://intranet.server.name: aport/>. Intranet.server.name is of course the Intranet server, and aport is the administrator port.

Your browser should come up to the administrator server page. Go to the line under Global URL Configuration that lists Server Port Number and change it from 80 to a different port number. After you

submit the form and restart the server it will be running at the new port.

Using Hard to Guess Names

Most companies want people to be able to find their Web sites so they usually name them www.company.com. This also makes it easy to remember.

However, there is nothing saying that the Web server has to be called www.company.com. It can be named something as obscure as udu33rf.company.com. There is not much of a chance of someone guessing a name like that.

Hiding Your Server's Name

Using a hard-to-guess name is useless if it is listed in the Domain Name Service (DNS), or shows up in public access logs. This section discusses how to setup your machine so it cannot be found easily.

NOTE

DNS is the system that allows machines on the Internet to know who they are talking to. It is a hierarchical, distributed naming system that allows each site to maintain its own list of hostname to IP address mappings. DNS stands for Domain Name Service.

Hiding Via Separate Name Servers

Some sites run separate internal and external name servers to allow internal machines to connect to the Intranet server by name, but not give out information to the rest of the Internet.

Basically it works like this. When an internal user wants to talk to a machine by name, queries go to the internal name server, which does one of two things. If it is a local machine, it replies with the correct answer. If it is an external machine, it queries the external name server to resolve it.

External users can only connect to the external server, which doesn't know about all the internal machines. Because it doesn't have internal names listed, it reports back an error whenever someone asks for them.

Setting up and maintaining two separate name server machines can be cumbersome and costly. Experts disagree on whether running separate name servers gains anything or not, but all experts agree that simply hiding your Intranet is not enough to protect it.

Keeping Your Name from Other Sites' Log Files

It is also important to keep your server name from appearing in other sites' access logs. Some sites don't adequately protect their access logs from being read from Internet users and this can allow your hard-to-guess name to be found out.

The best way to keep your server's name out of other sites' access logs is simply not to use it for Web surfing, Usenet posting, sending e-mail, or any other Internet access.

Whenever someone connects to a Web server it makes an entry in the log file. Some sites, either on purpose or by accident, make their log files publicly readable. If you use your secret Web server for browsing, it might get listed in one of these public log files and get indexed at a search engine site. All a hacker has to do is connect to a major search engine and search for yourcompany.com and look for strange names.

The same is true for Usenet postings and e-mail lists, which are usually archived on a Web site and indexed. It makes sense to occasionally search for your company name on search engines, to check for security problems such as this.

Pitfalls with Security Through Obscurity

Unfortunately there are many programs available that can search for open ports on a machine or a range of machines. These port scanners are available for downloading from the Internet and can run very quickly. In TCP/IP there is a limit of 65,535 different port numbers that can be used so it is not very hard for a hacker to scan all the ports on a host and find any hidden servers.

NOTE

Port scanning software is a hacker's friend, but it can also be a network manager's friend as well. Port scanning software can detect unauthorized servers running that can cause security problems. Port scanners are available from several places, including "strobe" from <ftp://suburbia.net/pub/> and "netcat" from <ftp://avion.org/src/hacks>.

It is possible to booby trap your machine to catch people running port scanning software. You can set up dummy servers that report when someone connected to them and where he or she connected from. Analyzing the logs from these servers may allow you to detect an attack before someone breaks in.

NOTE

TCP Wrappers (<ftp://ftp.win.tue.nl/pub/security>) are a set of programs written by Wietse Venema. They work by verifying and logging connection information before starting the server. For example, by setting wrappers around your httpd server you can detect what IP addresses have connected, if they really are who they say they are, and accept or refuse the connection. If it is accepted, then it connects with the Web server normally. (See figure 23.1.)

TCP Wrappers also can be used to set up dummy servers. They can be set up to send e-mail, send something to the printer, or send a page if someone tries to connect to a trapped port. Even if you don't need to set up wrappers around your servers, this notification feature makes TCP Wrappers worth downloading.

Figure 23.1 : TCP Wrappers checks the connection before allowing the server to talk to the client.

Using the Server Security

If simply hiding the site is not enough, what better defensive measures are available? Most WWW servers offer a layer of protection called access controls.

These access controls may be used to allow you to define IP address ranges that can retrieve documents from your Web server. Most Web servers also allow you to specify a username and password before allowing any documents to be retrieved.

These two ways of protecting the Web server from being abused and the problems with using them are covered in the next few sections. For specific server configurations see the chapter in Part II related to your Web server, or consult your Web server documentation.

There are two security models you can use to secure your Web server:

- All that is not allowed is denied. This means that you start out denying everyone access and only allowing certain machines access. This can usually be done by placing an * (or all) in the deny field and placing the local network number in the allow field.
- All that is not denied is allowed. This means that you allow everyone and deny those sites that are known to be bad. This would be done by placing an * (or all) in the allow field and listing out any site that shouldn't be allowed access to the deny field.

Security experts agree that it is best to deny everything and only allow in what is needed. This reduces the chance of the administrator overlooking something. In an Intranet it is better to deny everyone and allow local IP addresses, rather than try to list everyone who can't see what is on the Intranet.

CAUTION

Using the Web server security may stop a hacker from getting information from the Web server, but there are other ways to get information out of a machine, such as shared or exported drives. For more explanation, see the section on "Other Access Methods" later in this chapter.

Restricting by IP Address

Almost all Web servers have an access list that defines what machines or networks are allowed to retrieve documents or submit forms. This access list is usually made up of a list of allow and deny fields.

Using Apache, you can restrict access by IP address. This requires you to create a Limit directive in the access.cfg file. This file is usually in the httpd/conf directory.

The Limit directive should look like the following:

```
<Limit>
order deny, allow
Deny *
Allow 123.123.123. *
</Limit>
```

To deny access by IP address to the entire server using Netware Web server, you also need to edit the ACCESS.CFG file and add a Limit directive. The syntax will be the same as Apache uses. The ACCESS.CFG file is located in the \WEB\CONFIG directory.

Setting up access control for Netscape is done via the server manager page. Point your browser to the administration page. See the section on "Changing the Port with Netscape Server" for instructions. Go to the link under Access control labeled "Restrict access from certain addresses." This should bring up a new page. On this page you need to specify what resource you want to protect. To protect the entire server enter *. Then enter the IP addresses that are allowed access. This should be your local network number followed by an *. For example, if the local network is 123.123.123, then you would enter 123.123.123.* to allow all the hosts in your network to have access.

NOTE

TCP Wrappers also can be used to limit access to the Web server by IP address range. They can also be used to send an error message to external users notifying them of where the external server really is.

IP Spoofing

Restricting access based on the IP address of the requesting machine is a good start toward security; however, that security is only as good as the security of the IP protocol. The IP protocol was designed with ease of use, not security, and allows people to masquerade as other hosts. This is called IP Spoofing.

IP Spoofing is, in a nutshell, telling your machine to use someone else's IP address. This also requires some modifications to get the response to go back to the malicious site, but it can be done.

CAUTION

In some cases it isn't even necessary to get a response. For example, if a form automatically creates an account, then all a hacker would have to do would be to pretend to be a trusted host and send the URL to add an account to the Web server. Then simply wait for the account to be created. The hacker doesn't need to see a response; he just needs to be able to send commands.

IP Spoofing can be fixed using router access lists. All that is required is to tell the router not to allow any machines using your IP address in from the Internet. (See figure 23.2.) This way your Web server never sees the packets and doesn't have to worry about them. This won't work, though, if you are sharing your Web server with another company over the Internet because the hacker could pretend to be the other company.

Figure 23.2 : Routers can be used to prevent IP Spoofing attacks.

Usernames and Passwords

Using IP addresses as security protection is a good start but doesn't handle all cases. The most obvious case is when multiple people use the same machine but only one person should have access to your Intranet. This can happen if your users use Internet Service Providers (ISPs) to gain access when they are out of the office.

ISPs generally allow multiple people to log in to one machine or allow dial-up users to share an IP address. Either way it can be disastrous to a security policy based on IP addresses for access decisions.

In this case, it is required to use usernames and passwords. Most Web servers allow username security. (See Part II for specific server configuration information.) Setting up username security is covered in later sections.

When a user encounters a page that is protected, a box appears asking for username. After the username is entered, the password is required. Once the password has been entered, it is checked to make sure it is the correct one. If so, the document is sent to the user; otherwise, an error is sent back. Users usually only need to authenticate once per site, per session. Figure 23.3 shows a password prompt.

Figure 23.3 : Users are prompted for a username and password for secure pages.

One of the problems with username and password protection has to do with the HTTP protocol itself. HTTP sends text "in the clear," which means exactly as you type it. This would allow anyone with a network sniffer on a network segment your traffic goes through to see your username and password.

In Chapter 1 we discussed HTTP as well as the secure versions of HTTP (HTTPS and SSL). Using these protocols or an encrypted link will protect against sniffers, however not all servers support SSL and HTTPS. Check the server documentation in Part II before trying to use this security feature.

NOTE

Using the Web server security only limits what files can be retrieved through the HTTP protocol. Other ways to retrieve files must be protected as well. For example, a hacker may be able to use anonymous FTP to get to the protected documents.

Configuring Basic Authentication

Using usernames and passwords with Apache requires editing either the Directory directive in the access.cfg file or the file .htaccess file in the directory you want to protect. The Limit directive needs to contain the following:

```
<Limit>
require valid-user
AuthName [Auth-domain]
AuthType Basic
AuthUserFile [user-file]
AuthGroupFile [group-file]
</Limit>
```

The Limit directive is covered in more detail in [Chapter 3 "Installing and Configuring HTTPD for UNIX."](#)

Netscape's servers also allow username and password authentication. To enable it you will need to go into server manager and select the link under Access Control that reads, "Restrict access to part of your server through authentication." Follow the directions on this form. Once it is filled out and submitted you need to restart the server for changes to take effect.

Netware Web server allows either file-based authentication, like Apache, or NDS-based authentication.

To restrict server access, you must manually edit the global access.cfg file. Before the Limit directive for the DocumentRoot directory you must add the following three directives:

- AuthType Basic-This tells it what type of authorization type to use.
- AuthName [name]-This is a name telling the user which password to enter.
- AuthUserFile [filename]-This is the filename that contains the encrypted passwords.

Inside the Limit directive you need to add a Require directive telling the server which users can have access to the directory. This can be a list of usernames or "valid_user". Using valid_user allows any user listed in the AuthUserFile to have access as long as he or she enters the correct password.

The encrypted file is created using the command pwgen. It takes as arguments the input file and the output file. The input file is a file containing usernames and clear text passwords. For example:

```
Rich:secret
Tom:quiet
Al:ring
```

After running pwgen and copying the encrypted file in ServerRoot (usually SYS:WEB) you should copy the unencrypter file to floppy disk and remove it from the system. This is added protection to keep the passwords from being accidentally discovered.

If you choose to use NDS passwords, you can use the webmgr.exe program to configure the access restrictions. To do this, perform the following:

1. Start webmgr.
2. Click File/ Select Server.
3. Select the correct server directory (probably a drive mapped to SYS:\WEB).
4. Select the Directory from the drop-down list. To secure the entire server select \.
5. Select Directory Services from the Authorization list.
6. Type the NDS context name that contains the user object that should have access.
7. Select the user.
8. Click Add to Authorized users list.
9. Click OK.
10. Click Save and Restart.
11. Enter the Web server password and OK for your changes to take effect.

Microsoft IIS can also be used for username authentication. This can be done by opening the Web server properties box and selecting the service tab. Choose Basic under Password authentication. This will force users to supply a username and password. This username must be a valid account for the machine running IIS or in an accessible NT domain.

Other Access Methods

Earlier we mentioned other ways for hackers to get files without going through the Web server. This effectively gets around any security the Web server may be using.

Different operating systems have different ways to get files from the server, UNIX servers have FTP, NFS, TFTP, and others, while NT may have shared drives as well as FTP and TFTP. Netware users may also have FTP, TFTP, and shared drives. It is necessary to determine any way that a file can be retrieved from the server and protect against it.

TCP wrappers can be used to allow only certain IP addresses access to the servers that you know about, but there may always be a server you have missed. Hackers will, if possible, place servers running on high ports to allow themselves a way to get back in, just in case their initial break-in is discovered. Running port scanning software against your machines on a regular basis is a recommended practice to help catch these backdoor servers.

The other alternative is to place your Intranet server behind a firewall. This effectively blocks access to all services except the ones you allow. It is possible to use the other security models and disallow access to any servers that you know are bad, but you may miss some. It is always more secure to disallow everything and allow only what you know is needed.

In order to be sure all services are protected; the network must be told to refuse all connections and only allow through certain ones. This is a job for a firewall. Firewalls are covered in the next section.

Firewalls

Almost every company connecting to the Internet needs a firewall. The Web server software security is just not enough, because there are too many other ways to get information, such as FTP, TFTP, or shared file systems.

Firewalls are a system or group of systems that enforce a policy between two networks. In most cases one of the networks is the Internet; however, firewalls can be placed between any two networks.

CAUTION

Firewalls can protect against many types of attacks by limiting what sorts

of protocols can be passed into the company network. However, a firewall doesn't solve all the security problems. User education is required to teach users about security problems such as someone calling them on a phone and asking for a password. Dial-in modems are also a major concern.

Firewalls are split into two different categories: network-level and application-level firewalls. Both have their strengths and weaknesses and it is possible, and often desirable, to use a combination of the two type of firewalls, depending on your security requirements.

Network-level firewalls look at the packet header and see where it is trying to connect to, it then decides if it can pass or not. Application-level firewalls look at the packet and decide what the packet is trying to do and whether it is permitted or not.

Network-Level Firewalls

Network-level firewalls are devices or systems that look at the IP packet header and decide whether the address and port are allowed to pass through or not. It does not look at the data portion of the packet and does not understand what is going through.

The most common network-level firewall is simply a router with an access list defined on it. Machines with two network cards can also be used using special filtering software.

Network-level firewalls are very fast, because they only have to look at the IP header to decide whether the packet can pass or not. Network-level firewalls are also transparent. Because they don't interfere with the packet, users don't need to learn any special commands.

NOTE

Because network-level firewalls pass traffic directly between the two networks, both networks must have valid Internet address ranges.

Network-level firewalls however have shortcomings. The logging they use is often not as sophisticated as application-layer firewalls. Network-level firewalls also don't understand what they are passing and can only refuse or deny the protocol. For more sophisticated logging, or finer access control, application-level firewalls are needed. They are discussed later in this chapter.

Benefits of a network-level firewall are:

- Easy to configure
- Transparent to users
- Very fast
- Cost effective

Disadvantages of a network-level firewall are:

- Access control is based only on address/ports
- Logging is fairly simple
- Must have a valid address range

Router Access Lists

Many routers allow the administrator to define an access policy. These access policies tell the router what

packets are allowed to pass and which ones are to be dropped.

A simple access policy would say:

- Rule 1: Allow port 25 (e-mail) from anyone into the mailhost.
- Rule 2: Allow port 80 (WWW) from anyone into the external Web server.
- Rule 3: Allow users from 123.234.123 network log in to port 23 (telnet) to mailhost.
- Rule 4: Don't allow anything else in.

This simple access list allows anyone to send e-mail to the mailhost (Rule 1), and allows customers to get to the external WWW server (Rule 2). It also allows people on the network 123.234.123 to log in to the machine mailhost. The most important line is the last one, which says no one else can get in.

Access lists can and will be much more complex than this. Some routers allow you to define what to do if a packet is refused, such as log the attempt to a host or printer. You can also use other keywords to allow traffic out from the internal network to the rest of the world.

Access lists generally can be placed on the external port, on the internal port, or both, depending on the router.

Allowing traffic to a specific set of machines is commonly referred to as a screened-host firewall. This means that traffic is allowed to a specific host or set of hosts but only if it fits through the screen. Our access policy is an example of a screened-host firewall. (See figure 23.4.) The machines that external machines are allowed to talk to are called *bastion hosts*. Bastion hosts are hosts that are secured to (hopefully) resist attack.

Figure 23.4 : Screened host firewalls allow certain types of traffic to a bastion host.

In addition to the screened-host firewall is a screened-subnet firewall. This allows a defined set of traffic through to an entire network. These configurations are covered in more detail later in the chapter.

Packet Filters

Packet filters are usually machines with two network cards running special software to allow or disallow packets based on their address and destination port.

Packet filters are similar to routers and can be used to build screened-host or subnet firewalls. Packet filtering software can be downloaded from the following sites:

- IP Filter-<ftp://coombs.anu.edu.au/pub/net/kernel>
- KarlBridge-<ftp://ftp.net.ohio-state.edu/pub/kbridge>
- Screend-<ftp://decuac.dec.com/pub/sources>

Packet filter software is available for both DOS machines and UNIX machines and supports many different network cards. Packet filters are often used to act as firewalls because they can run on very cheap hardware. A 386 processor-based machine has more than enough CPU power for a 56 Kbps line to the Internet and can handle T1 speeds, depending on the exact configuration.

Application-Level Firewalls

Application-level firewalls are very secure hosts with two network cards: one set up on the internal side and one setup to run on the hostile side. This bastion host runs proxy servers for each protocol that needs to get through the firewall. This allows no traffic to pass between the two networks and the proxy must instead

interpret and relay data back and forth.

Application-level firewalls must understand the protocol they are proxying for and thus can do very sophisticated logging and access controls. For example, a proxy for HTTP may be set up to allow getting pages but not images. It could also be configured to disallow running a CGI program as in figure 23.5.

Figure 23.5 : Application level gateways can be configured to disallow certain parts of a protocol. Here we allow GETting pages but not POSTing scripts.

Benefits of an application-level firewall are:

- Very fine access control
- Better logging
- Can be used with reserved or unregistered IP address ranges

The disadvantages of application-level firewalls are:

- Harder to implement
- Performance isn't as good as with a network level firewall
- Special software required for each protocol
- Not always transparent to the user

Address Translation

Because application-level firewalls relay messages instead of passing traffic through, they can be used as Network Address Translators (NATs). NATs are useful for IP networks that are using unregistered address ranges or reserved ranges.

NOTE

With the explosive growth of the Internet it was realized that there were not enough addresses for everyone who used IP to have a registered address. As a result, RFC 1597 was written. It advised setting aside different ranges of addresses not to be used on the Internet, but to be used only on internal networks. Of course some people who started out using reserved numbers decided to get on the Internet and now must use NATs to perform address translation for them.

Because the Internet is designed not to route traffic to any reserved IP address ranges, using them for internal networks can help to prevent attacks from the Internet. This added security is not without a cost, though, because every external connection may require a network address translator to work properly.

Proxy Gateways

Most popular applications have proxy servers written for them. For example, Trusted Information Systems (TIS) has written a firewall toolkit that contains proxy servers for rlogin, telnet, FTP, X-windows, HTTP, and NNTP (Usenet). These proxies understand the protocol and can be used to safely allow these applications through the firewall.

Whether or not you choose to allow these protocols into your network requires looking at the security policy and deciding if they are required. Just because a protocol can be safely allowed through the firewall doesn't mean it should be allowed into the networks.

For a proxy to work it must understand the protocol it is trying to proxy. This means new protocols may not have a proxy written for them yet. These protocols either need to be passed through a network-level firewall or a generic proxy, also known as application forwarders.

Application Forwarders

Application forwarders are sort of emergency measures for firewall administrators. They allow proxy-like support for protocols that do not have special proxy software written for them yet. Application forwarders are also called generic proxies.

TIS includes a software package called plug-gw, which allows any protocol to be forwarded through an application level firewall without having to use a special proxy package.

Plug-gw may not allow as fine control as a special proxy, because it does not necessarily know what the protocol is doing. It can, however, log connection attempts, times, and hosts. It is a sort of compromise between network level firewalls and full proxy support for a protocol.

Socks

Socks is a generic application forwarder that is often used to allow access out to the world. Using Socks allows IP connectivity to hosts that are normally blocked by a firewall. Socks basically accepts connections from the internal network and relays these connections to the Internet hosts. It also relays data back and forth after the connection has been allowed and established.

Socks comes in source code form and must be compiled for each server platform. Socks is fairly easy to compile and can be done in a short amount of time. You can download socks from <ftp://ftp.nec.com/pub/security/socks.csrc>.

Socks is supported for SunOS, Solaris, AIX, SCO, and Linux.

Firewall Configurations

Firewalls are configured in different ways, depending on how your company decides to balance the cost versus security issue. The most secure firewalls often use a combination of security layers to offer the most protection.

There are three types of basic firewall configurations in use today. They are:

- Dual-homed gateways
- Screened-host gateways
- Screened-subnet gateways

Each configuration has a different tradeoff among between security, cost, and performance. For example, dual-homed gateways are easier to configure and set up than screened hosts, but at a slight loss in security.

Dual-Homed Gateways

Dual-homed gateways are hosts configured to support two network interfaces. One faces in toward the secure network and one faces out toward the hostile network. They are configured not to route traffic between them. (See figure 23.6.)

Figure 23.6 : Dual-homed gateways are built with two interfaces-one on each network. They don't allow routing.

The dual-homed machine is also called a bastion host. This machine needs to be as secure as possible, because it is the way a hacker will try to get in from the Internet.

The bastion host is the host that would normally run the external WWW server and FTP server as well as act as a mailhost for the Internet. It also will be the host that runs proxy servers for the users to get access to the Internet.

Dual-homed gateways are the cheapest of the three basic firewalls. However, they have one disadvantage over the other types of firewalls. They have a single point of failure. If any piece of software allows a hacker in, the entire network is exposed.

Screened-Host Gateways

Screened-host gateways are built using a screening router to block traffic to the internal network and to allow traffic to a bastion host. This host has one network interface as opposed to the dual-homed gateway that has two. (See figure 23.7.)

Figure 23.7 : Screened-host gateways are made up of a router and a bastion host.

Because most Internet service providers provide a router at the site, it is easy and often economical to use this configuration. Screened hosts have the advantage of being able to allow certain applications in to the network easily.

The ability to poke holes in a firewall can be a major disadvantage if the firewall administrator is forced to open holes for unsafe protocols.

Another disadvantage to screened-host gateways is the fact that this system relies on two separate security devices: the router and bastion host. If either of these fail, the network is exposed.

Screened-Subnet Gateways

A screened subnet is made up of two screening routers with the bastion in the middle. This makes a small isolated network between the secure and hostile networks. The bastion host sits in this isolated network. (See figure 23.8.)

Figure 23.8 : Screened-subnet gateways have an isolated network separated from both networks via routers.

It is possible to have multiple hosts in this isolated network, also called Demilitarized Zone or DMZ. This can help to alleviate performance problems.

This configuration is the most expensive but it is also the most secure because two devices need to fail to allow the network to be exposed.

Like the screened-host gateway, it is possible to poke holes through both routers to allow protocols in. This should only be done if the protocol is safe and there is no proxy available for it.

General Security

Many times network administrators put a firewall in place and assume that they are safe. However, this is very rarely the case. This is similar to putting a large steel door on the front of your house, and not paying attention to the glass windows scattered around the building.

Modems

Many companies have modems on their employees' machines. These can be for transferring files, connecting to online services, or might not be used at all. Many of these modems might be set to Auto Answer when called.

Auto Answer modems by themselves aren't a problem, especially on PCs that usually don't allow logins. However, UNIX machines are normally configured to allow remote logins on modem ports. This allows people in without going through the firewall.

UNIX machines aren't the only ones with remote capabilities, though. Many PC modem packages allow a remote user to take over a PC just as if he were sitting at the keyboard. This includes any network access that the machine may have.

Even worse than a login session is a PPP or SLIP connection. Many newer systems, such as Windows 95 or Windows NT, come with PPP software built in. It is possible for a user, knowingly or not, to enable this access. This allows the hacker to access any machine on the network virtually untraced.

It is important for a security administrator to be aware of these potential security holes. Hackers have access to many phone number scanners. These scanners work much like port scanners or network scanners. They start dialing numbers until they get a modem tone and then try to break in. Knowing which phone lines are connected to modems and securing them is as important as configuring a firewall.

Other Remote Networks

In addition to the Internet link, your company may have other remote networks that can be a source of attack. These may be something like Tymnet or Telenet, or a dedicated or dial-up link to another company.

Regardless of where these links go, you should always assume they connect to a hostile network and security should be in place on them. This may include a firewall or another safety device.

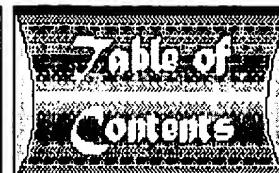
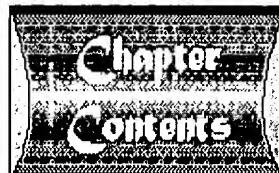
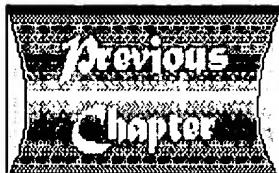
Physical Security

Physical security is usually not overlooked because this is the part of security most people are familiar with. Still some security teams spend more time looking for large-sized items rather than high importance items.

Many companies use 8mm or 4mm tape drives for backups. These tapes can easily fit in a shirt pocket. This is one way for data to leak out right past the firewall. Of course, employees need to be trusted to a point in any company.

Another problem can be unshredded confidential documents. One effective trick of a system hacker is to go trashing, which is basically going through a company's dumpster for information that may allow a hacker to get in. People often jot down usernames and passwords as well as telephone numbers and other security information. These notes often end up in the trash after they have been used. Shredding documents is a cost effective way of eliminating the usefulness of trashing.

If users and administrators understand the importance of securing the Intranet, many of these problems are taken care of. Having a secure Intranet can make the difference between people using the Intranet for work or for a toy. Making sure people realize they are accountable for the amount of time they spend on the Internet can also help to keep the Intranet from being abused.



EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	42116	port near3 (forward\$3 sourc\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L2	4152	(port near3 (forward\$3 sourc\$3)) same ((first 1st primary default) near3 port) same ((second\$3 2nd) near3 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L3	3054	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L4	95	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L5	81	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https) and network	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L6	81	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L7	768	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L8	641	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L9	21	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) same (prox\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L10	22	(port near3 (forward\$3 sourc\$3 redirect\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) same (prox\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L11	28	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and ((first 1st primary default) near2 port) same ((second\$3 2nd) near2 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L12	40	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and ((first 1st primary default) near2 port) and ((second\$3 2nd) near2 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L13	12	L12 not L11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L14	201	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and (software code program)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L15	201	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and (software code program) and (port forward\$3 sourc\$3 redirect\$3 prox\$3 first 1st primary default port second\$3 2nd port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L16	0	713/201.ccis.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L17	299	"5623601"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L18	75	"5623601" and ssl	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L19	70	"5623601" and ssl and https	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L20	16	08/322078	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L21	2816	http with (encrypt\$3 encapsulat\$3 encipher\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L22	1624	http near5 (encrypt\$3 encapsulat\$3 encipher\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L23	19648	(http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L24	21751	(http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5 forward\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L25	309	@ay < "1997" and (http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5 or forward\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L26	137	@ay < "1997" and (http same (port near\$3 (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L27	0	@ay < "1997" and (http same (port near3 (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L28	1	@ay < "1997" and (http same (port with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L29	127	(http same (port with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L30	127	(http port redirect\$3 "80" "443" https forward\$3 swap\$4) and L29	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L31	20	(http same (url with (redirect\$3 with port)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L32	692	(http same (url with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L33	38	(http same (url with (redirect\$3))) same port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L34	55	(http and (url with (redirect\$3))) same port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L35	653	(http and (url with (redirect\$3))) and port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L36	105	(http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L37	14721417	@ay < "1997" (http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L38	1	@ay < "1997" and (http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L39	5	@ay < "1997" and (http and (url with (\$2direct\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L40	11	https near3 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L41	13	https near5 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L42	13	"https" near5 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L43	14	"https" near7 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L44	0	secure with http near7 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L45	0	secure with http with :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L46	86	secure with http with "80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L47	26	ssh same "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L48	2	ssh with (forward\$3 redirect\$3) with "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L49	4	ssl with (forward\$3 redirect\$3) with "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L50	794	(protocol near9 (forward\$3 redirect\$3)) same http same (ssl "https" ssh)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L51	16	(protocol near9 (forward\$3 redirect\$3)) same http same (ssl "https" ssh) and @ay <"1997"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L52	81	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https) and network	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L53	768	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L54	0	@ay < "1997" and (http same (port near3 (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L55	40	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and ((first 1st primary default) near2 port) and ((second\$3 2nd) near2 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L56	299	"5623601"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L57	75	"5623601" and ssl	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L58	3054	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L59	70	"5623601" and ssl and https	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L60	19648	(http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L61	2816	http with (encrypt\$3 encapsulat\$3 encipher\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L62	42116	port near3 (forward\$3 sourc\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L63	4152	(port near3 (forward\$3 sourc\$3)) same ((first 1st primary default) near3 port) same ((second\$3 2nd) near3 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L64	21751	(http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5 forward\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L65	0	713/201.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L66	653	(http and (url with (redirect\$3))) and port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L67	127	(http same (port with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L68	692	(http same (url with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L69	14721417	@ay < "1997" (http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L70	0	secure with http near7 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L71	0	secure with http with :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L72	794	(protocol near9 (forward\$3 redirect\$3)) same http same (ssl "https" ssh)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L73	1	@ay < "1997" and (http same (port with (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L74	95	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L75	1	@ay < "1997" and (http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L76	81	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near3 port) with ((second\$3 2nd) near3 port) and (ssl https) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L77	21	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) same (prox\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L78	22	(port near3 (forward\$3 sourc\$3 redirect\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) same (prox\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L79	28	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and ((first 1st primary default) near2 port) same ((second\$3 2nd) near2 port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L80	12	L55 not L79	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L81	127	(http port redirect\$3 "80" "443" https forward\$3 swap\$4) and L67	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L82	20	(http same (url with (redirect\$3 with port)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L83	16	08/322078	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L84	38	(http same (url with (redirect\$3))) same port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L85	55	(http and (url with (redirect\$3))) same port	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L86	5	@ay < "1997" and (http and (url with (\$2direct\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L87	11	https near3 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L88	105	(http and (url with (redirect\$3))) and (port with (forward\$3 swap\$4 redirect\$3))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L89	13	https near5 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L90	13	"https" near5 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L91	14	"https" near7 :80	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L92	26	ssh same "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L93	137	@ay < "1997" and (http same (port near\$3 (redirect\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L94	2	ssh with (forward\$3 redirect\$3) with "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L95	16	(protocol near9 (forward\$3 redirect\$3)) same http same (ssl "https" ssh) and @ay <"1997"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L96	201	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and (software code program)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L97	86	secure with http with "80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L98	4	ssl with (forward\$3 redirect\$3) with "port 80"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44

EAST Search History

L99	201	(port near3 (forward\$3 sourc\$3 redirect\$3)) with (prox\$3) and (software code program) and (port forward\$3 sourc\$3 redirect\$3 prox\$3 first 1st primary default port second\$3 2nd port)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L100	309	@ay < "1997" and (http same (port near\$3 (exchang\$5 or redirect\$5 or swap\$5 or chang\$5 or switch\$5 or forward\$3)))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L101	641	(port near3 (forward\$3 sourc\$3)) with ((first 1st primary default) near2 port) with ((second\$3 2nd) near2 port) and network\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44
L102	1624	http near5 (encrypt\$3 encapsulat\$3 encipher\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/12/11 09:44